

## SYSTEM AND METHOD FOR NETWORK USER INTERFACE OLAP REPORT FORMATTING

### Related Applications

This application is related by subject matter to co-pending applications entitled "System and Method For OLAP Report Generation With Spreadsheet Report Within the Network User Interface," Attorney Docket No. 53470.000006 and "System and Method for Asynchronous Control of Report Generation Using a Network Interface" Attorney Docket No. 53470.000005, both of which are filed herewith.

### Field of the Invention

This invention relates to a system and method for generating, formatting, and styling OLAP reports in a network interface.

### Background of the Invention

The ability to act quickly and decisively in today's increasingly competitive marketplace is critical to the success of any organization. The volume of data that is available to organizations is rapidly increasing and frequently overwhelming. The availability of large volumes of data presents various challenges. One challenge is to avoid inundating an individual with unnecessary information. Another challenge it to ensure all relevant information is available in a timely manner.

One known approach to addressing these and other challenges is known as data warehousing. Data warehouses and data marts are becoming important elements of many information delivery systems because they provide a central location where a reconciled version of data extracted from a wide variety of operational systems may be stored. As used herein, a data warehouse should be understood to be an

informational database that stores shareable data from one or more operational databases of record, such as one or more transaction-based database systems. A data warehouse typically allows users to tap into a business's vast store of operational data to track and respond to business trends that facilitate forecasting and planning efforts.

5 A data mart may be considered to be a type of data warehouse that focuses on a particular business segment.

Decision support systems have been developed to efficiently retrieve selected information from data warehouses. One type of decision support system is known as an on-line analytical processing system ("OLAP"). In general, OLAP systems analyze  
10 the data from a number of different perspectives and support complex analyses against large input data sets.

There are at least three different types of OLAP architectures - ROLAP, MOLAP, and HOLAP. ROLAP ("Relational On-Line Analytical Processing") systems are systems that use a dynamic server connected to a relational database  
15 system. Multidimensional OLAP ("MOLAP") utilizes a proprietary multidimensional database ("MDDB") to provide OLAP analyses. The main premise of this architecture is that data must be stored multidimensionally to be viewed multidimensionally. A HOLAP ("Hybrid On-Line Analytical Processing") system is a hybrid of these two.

20 ROLAP is a multi-tier, client/server architecture comprising a presentation tier, an application logic tier and a relational database tier. The relational database tier stores data and connects to the application logic tier. The application logic tier comprises a ROLAP engine that executes multidimensional reports from multiple end

users. The ROLAP engine integrates with a variety of presentation layers, through which users perform OLAP analyses. The presentation layers enable users to provide requests to the ROLAP engine. The premise of ROLAP is that OLAP capabilities are best provided directly against a relational database, *e.g.*, the data warehouse.

5 In a ROLAP system, data from transaction-processing systems is loaded into a defined data model in the data warehouse. Database routines are run to aggregate the data, if required by the data model. Indices are then created to optimize query access times. End users submit multidimensional analyses to the ROLAP engine, which then dynamically transforms the requests into SQL execution plans. The SQL is submitted  
10 to the relational database for processing, the relational query results are cross-tabulated, and a multidimensional result set is returned to the end user. ROLAP is a fully dynamic architecture capable of utilizing precalculated results when they are available, or dynamically generating results from atomic information when necessary.

The ROLAP architecture directly accesses data from data warehouses, and  
15 therefore supports optimization techniques to meet batch window requirements and to provide fast response times. These optimization techniques typically include application-level table partitioning, aggregate inferencing, denormalization of data structures, multiple fact table joins, and the use of specific RDB optimizer tactics which vary with each particular brand of relational database.

20 MOLAP is a two-tier, client/server architecture. In this architecture, the MDDB serves as both the database layer and the application logic layer. In the database layer, the MDDB system is responsible for all data storage, access, and retrieval processes. In the application logic layer, the MDDB is responsible for the

execution of all OLAP requests. The presentation layer integrates with the application logic layer and provides an interface through which the end users view and request OLAP analyses. The client/server architecture allows multiple users to access the multidimensional database.

5 Information from a variety of transaction-processing systems is loaded into the MDDB System through a series of batch routines. Once this atomic data has been loaded into the MDDB, the general approach is to perform a series of batch calculations to aggregate along the orthogonal dimensions and fill the MDDB array structures. For example, revenue figures for all of the stores in a state would be added  
10 together to fill the state level cells in the database. After the array structure in the database has been filled, indices are created and hashing algorithms are used to improve query access times.

Once this compilation process has been completed, the MDDB is ready for use. Users request OLAP reports through the presentation layer, and the application  
15 logic layer of the MDDB retrieves the stored data.

The MOLAP architecture is a compilation-intensive architecture. It principally reads the precompiled data, and has limited capabilities to dynamically create aggregations or to calculate business metrics that have not been precalculated and stored.

20 The hybrid OLAP ("HOLAP") solution is a mix of MOLAP and relational architectures that support inquiries against summary and transaction data in an integrated fashion. The HOLAP approach enables a user to perform multidimensional analysis on data in the MDDB. However, if the user reaches the bottom of the

multidimensional hierarchy and requires more detailed data, the HOLAP engine generates an SQL statement to retrieve the detailed data from the source relational database management system ("RDBMS") and returns it to the end user. HOLAP implementations rely on simple SQL statements to pull large quantities of data into  
5 the mid-tier, multidimensional engine for processing. This constrains the range of inquiry and returns large, unrefined result sets that can overwhelm networks with limited bandwidth.

As described above, each of these types of OLAP systems are typically client-server systems. The OLAP engine resides on the server side and a module is typically  
10 provided at a client-side to enable users to input queries and report requests to the OLAP engine. Current client-side modules are stand alone software modules that are loaded on client-side computer systems. One drawback of such systems is that a user must learn how to operate the client-side software module in order to initiate queries and generate reports.

15 Although various user interfaces have been developed to enable users to access the content of data warehouses through server systems, many such systems experience significant drawbacks. One system enables the user to access information using a network interface device (*e.g.*, a web browser). Such systems provide the advantage of allowing users to access the content of a data warehouse through existing web  
20 browser applications residing on their desktop or other personal computer connected over a network to the server system. Such systems however suffer from various drawbacks.

In existing web-based user interfaces, a user inputs a query or report request through the web browser, the web browser initiates the report to the server, and waits until the report finishes before providing results back through the web browser to a user. While the web browser is waiting for the report, the web browser "locks up,"

5 thereby requiring that the user open another browser window if another site or page is desired to be opened. Accordingly, if the delay between the report initiation and completion is lengthy, some users may incorrectly assume that the delay is caused by an error, when in reality the delay may simply be caused by the complicated nature of the query. Users may become impatient and resubmit the report by selecting the

10 reload or refresh feature of the web browser. Current systems treat each such request separately and, therefore, each time the user resubmits a request, another query for the same information is submitted to the server. Because the reason for the delay in most circumstances is the size or complexity of the query, the server is then hit with processing multiple complex or large queries simultaneously when only one such

15 query or report is desired. Accordingly, these multiple reports tax server resources unnecessarily.

Another drawback this example illustrates is that current systems focus on requests. Accordingly, each request is treated as a different request that is submitted to the server for processing. Accordingly, identical requests initiated by different

20 users are submitted to the server system for processing, thus utilizing the server resources twice to generate the same report. These and other drawbacks exist with current web-based data warehouse/server system user interfaces.

### Summary of the Invention

An object of the invention is to overcome these and other drawbacks in existing systems.

Another object of the invention is to provide a system and method for enabling  
5 formatting of OLAP reports for presentation in a network user interface.

Another object of the present invention is to provide a system and method that enables selected of predefined auto-formats that may be applied to a selected report or series of reports.

Another object of the invention is to provide a system and method for  
10 providing levels of formatting to be applied to multiple OLAP report outputs for presentation in a network user interface.

These and other objects of the present invention are realized according to various embodiments of the present invention. Accordingly, an embodiment of the present invention comprises a network-based system for enabling users connected  
15 over the network to an OLAP system to select formatting options for reports requested for processing by the OLAP system. The system enables a user to select a workbook comprising one or more reports to be requested for processing by the OLAP system. The user may also select a single stand-alone report. The system then enables a user to select the format for each report in the workbook, for one or more reports in the  
20 workbook or to a stand-alone report. The format specified may be multiple levels of format that are applied in a hierarchical manner including global formats that apply to all reports in a workbook, report formats which apply to specific reports in the workbooks, attribute/element/metric formats which apply to specific attributes,

elements, and metrics of a report, and value formats which apply to specific values within an element of a report. The formats specified may also comprise merged report formats that specify the merger of two or more reports, combined grid and graphs within a report and an autoformat that enables each recipient of the report to specify  
5 the format in which that recipient views the report when the report is received.

Other objects and advantages of the present invention exist.

### **Brief Description of the Drawings**

Fig. 1 is a schematic block diagram of a method for generating reports asynchronously in accordance with an embodiment of the invention.

10 Fig. 2 is a schematic block diagram of a method for generating, monitoring, and canceling reports asynchronously in accordance with another embodiment of the present invention.

Fig. 3 is a schematic illustration of a report list in accordance with one embodiment of the invention.

15 Fig. 4 is a schematic illustration of a status page in accordance with an embodiment of the invention.

Fig. 5 is a schematic block diagram of an asynchronous report requesting network system in accordance with an embodiment of the invention.

20 Fig. 6 is a schematic block diagram of system components in accordance with an embodiment of the invention.

Fig. 7 is a schematic block diagram of a network-based system in which the asynchronous report generation system may operate in accordance with one embodiment of the invention.



### **Detailed Description of the Preferred Embodiments**

Accordingly, a method of operating a network-based OLAP system is provided. The system may comprise an Internet-based web browser interface that connects users over the Internet to a web server system that connects to the OLAP  
5 server system for processing requested reports.

A method 200 may be provided according to an embodiment of the present invention, as depicted in Fig. 1. In this method, a user system providing a user interface module may be connected over a network to an OLAP server system. The first step in method 200 then involves step 202 whereby the user may request the  
10 processing of a report to be run on the OLAP system through a web browser or other user interface system. The report may then be transmitted by the web browser or other user interface system to the server system such as a network module in communication with the web server.

After a report has been requested, in step 204, the system compares that report  
15 with other pending or recently completed reports for the user submitting the report and for other users connected to the system to determine if the requested report is the same as a pending or processed report. In an embodiment, if a particular report has recently been completed or is already in the queue to be processed by the server system, the system does not create a new report to be processed. Rather, in step 207, the system  
20 adds the user to the report that has already been requested so that when the previously requested report completes, then all users that requested that report receive the results. If the report has already completed, the system immediately returns the results to the

user as part of steps 215 and 226. If the report has not been requested, then in step 208, a report request is generated and transmitted to the OLAP processing system.

In step 226, the system returns control of the network interface, such as a web browser or other user interface module, to the user so that the user may perform other tasks. Specifically, the server system may direct the web server to transmit a signal to enable the user to use the web browser to perform other tasks, including submission of another report. The web server may also transmit a page that informs the user that the report that was requested was received and is scheduled for processing, such as a status page as discussed below.

Then, in step 220, the system notifies the user when a report has been completed so that the user may review the results of the report. By returning control of the web browser to the user upon receipt of the report request, the system frees the user to submit multiple reports for processing in an asynchronous manner. Additional details and features of the present invention may be understood with respect to Figs. 2-7 below.

Fig. 2 illustrates a more detailed embodiment of method 200 in accordance with one embodiment of the invention. Specifically, in step 204, the system determines whether that report is the same, substantially the same or similar to a report that has previously been requested. To do so, the system may store or have access to all, or a certain predetermined subset of all, reports forwarded to the system within a predetermined period of time (e.g., several hours, several days, or several weeks).

In an embodiment, the system, in step 204, may compare the name of the report with the name of other reports maintained in a cache that comprise all reports already pending and reports that have been completed but not removed from the cache. Also, the system may compare the template, filter, and user view of a request to the template/filter/user view combination of other requested reports in the cache. If the template/filter/user view combination selected are identical, then the reports may be deemed to be the same.

Also, two reports may also be determined to be the same for purposes of avoiding generating a new report even if the templates or filters are not identical. For example, two reports may be deemed to be substantially the same if the template/filter/user view combination is within a predetermined range of a template/filter combination of a previously run or pending report. For example, if a report for gross sales of running shoes for March, 1999 in the U.S. has been forwarded for processing, it may be determined that a request for gross sales of a particular brand of shoes for March, 1999 in the U.S. is substantially the same or similar enough to be the same report. Other methods of determining similarity of reports to reduce the number of jobs to be processed by the system may also be used.

In step 208, if the report is different from other pending or recently completed reports, then the report is generated by placing the report in the processing queue in a known manner. Additionally, the report may be added to a report list in step 210. The report list may comprise a list of all reports generated within a predetermined period of time and may include pending and completed reports. The report list may contain a report entry for each report containing the names or identifier(s) of system users that

have requested that report. The report entry for each report may also contain additional information if desired (*e.g.*, complete request parameters, user id, requests GUID, electronic mail address, employee numbers, metadata user id, project name, report types (prompt or non-prompt), auto-prompt selections, report description, report name, filter name, template name, folder name, requested time, submitted time, completed time, filter details, status, and operators being applied (such as AND, OR, and NOT)). Whether the report is new or previously requested, the request is then added to the report entry for the report in step 212. The request may include user name and/or other identification, grid/graph selected, time requested, etc. to identify the request.

It may also be desirable to notify the user of the status of the report request so in step 214, a status page may be generated. Therefore, in step 214, a status page may be generated and presented to the user by the system over the network. The status page may identify the name of the report, date and time submitted, and/or any other additional information. An example of a status page is depicted in Fig. 4. The status page may be displayed as a separate window and may be accessed by the user through viewing a report list for that user as described below. At this point, the system may enable the user to request another report, view old reports or perform other types of actions.

The status page may also enable the user to cancel a particular request, as for example, illustrated in Fig. 4 and described below. The user may also be able to cancel a report from a user report list, as for example, the user report list shown in Fig.

3, such as by clicking on the report entry and selecting a cancel icon, pull-down menu option, or other mechanism for canceling a report.

By treating each report separately and asynchronously, a report may be canceled at any time. As each report may have multiple users that have been added to that report, however, in one embodiment, a report is not canceled from the queue at the system until all users that have been added to that report have canceled. If a user cancels the report, the user's name or id is removed from the report entry for that report and that canceled user then cannot access the report results unless they resubmit another request for the report. The details of this process are depicted in steps 216-226.

Specifically, in step 216, the system determines whether a report request has been canceled. If a report request has been canceled by a user, the system removes the request for that from the report list in step 218. After removing the user from the report entry for that report, the system, in step 220, determines whether any users remain in the report entry for that canceled report. If any user names remain, method 200 returns to step 216. If no other user names remain in the report entry, the system terminates the report and deletes it from the queue of reports being processed by the server system in step 222.

If, in step 216, a report request has not been canceled, the system determines whether any of the reports are complete in step 224. If a report is complete, in step 226 the system notifies all of the system users identified in the report list for that report that the report has been completed. Notification may be made using e-mail, facsimile, pager, through the web interface such as through the web browser or other

user interface module using push technology, etc. or any other mechanism as described below.

If the report generation is not complete, in step 224 the system returns to step 216 to determine whether any report request cancellations have been entered into the system. This continues until the report generation is complete and all systems users are notified that the report generation is complete for all reports on the system. The process may involve polling whereby steps 216 and 224 are performed on a predetermined schedule, such as every seven seconds or so. Other schedules may also be used.

Method 200 provides numerous advantages for generating reports. Receiving report requests and processing those requests asynchronously enables users to operate other applications while a report is being generated. Additionally, by focusing on the identity of the report rather than a request for that report, duplicative report generation is minimized. Therefore, if several users request the same report, only one report is generated.

Fig. 3 illustrates a user-specific report list 300 in accordance with an embodiment of the invention. Report list 300 may contain a report list for reports submitted for a specific user within a predetermined period of time. That time may be a particular log-in session or a discrete time interval for example. For each report, report list 300 may contain a "name" field 302 which identifies the name of the report being generated. A folder field 303 may be provided for indicating the folder or other category that the report may be organized. A "submitted" field 304 may be used to identify the date and/or time the report request was submitted. Other fields, for

example, a “details” field 306 may be used to provide additional information, such as, the status of the report. Fields 302-306 may be sorted according to various criteria (*e.g.*, ascending or descending order) by, for example, clicking on a field using a conventional computer mouse. It is to be understood that report list 300 may be  
5 sorted according to the value of a field (*e.g.*, user, date, report, etc.). A status field may also be provided and contain information as described below with respect to Fig. 4. Further, various icons or other symbols may be used to indicate the status or type of report.

As Fig. 3 depicts, various types of reports may be processed according to the  
10 present invention and these different report types may be displayed differently in report list 300. For example, an auto-prompt report may be provided with details of the prompting to show the filters, metrics, what information has been drilled on, what was selected in the prompt, other filter details, and other parameters selected for the report.

15 As discussed above, to notify the user that the request has been received and is being processed, a status page may be presented to the user. A status page may be used to display the progress of a report being generated and a user-specific report list may be viewed.

Fig. 4 illustrates a status page 400 in accordance with one embodiment of the  
20 invention. Status page 400 provides a view displaying the current status of a report’s execution cycle. The execution cycle identifies the progress of a report generation (*e.g.*, 10% complete, Error, complete with errors, pending, complete, sending request, generating SQL, submitting job request, request in queue, running report, retrieving

results, cross-tabbing data, formatting data, report complete, error running report, report canceled, m of n reports completed). Status page 400 may contain a graphical display 402 of the progress of the report generation on a percentage complete basis. Selectable icons may also be displayed within status page 400. For example, a report  
5 list icon 404 may be displayed which enables a user to display a report list. A "refresh" icon 406 may be used to manually update the progress of the report generation. As described in further detail below, this operation may be automated and set by a user. Additionally, a "cancel request" icon may be displayed which permits a user to cancel the generation of a requested report.

10 A refresh status page rate may be set by a system user or by a system administrator. The system maintains a master list of requested reports and stores the list on a server. The master list may be refreshed when caches expire, a report is canceled, or a report is removed from a cache for another reason, etc.

A system administrator master report list may also be provided according to  
15 the present invention and enabled through network output module 22. The administrator report list may be similar to user report list 300 but may additionally comprise entries for each report to identify the users that have requested that report. Further, the system administrator report list may comprise all reports run within a predetermined period of time and may be broken out project by project if desired. The  
20 report list may provide a mechanism that enables a system administrator to delete entire reports, groups of reports, to delete a particular user or groups of users from a particular report or modify the order of reports in the queue. The administrator may thereby control reports submitted to the server system. In one embodiment, all such



reports may be stored in a cache and the master report list may contain all cached reports plus all pending reports.

The information in the history list may include the name of the request, the report display mode (graph or grid), the date and time the request was last submitted, the status of the request, the users who have submitted the request and their identifiers or keys, filter details for report requests, and workbook report names for workbook requests. The administrator may refresh the list, delete a request or view details about a request. The details may include the filters and templates requested.

According to another embodiment of the present invention, users may be able to submit requests for workbooks to be processed by the server system on data in the data warehouse. A workbook may comprise complex reports that generally comprise a plurality of reports and template/filter combinations that are collected to be generated as a large complex report. Therefore, as the term "report" as used herein should be understood to comprise a single template/filter combination run against data in the data warehouse or a collection of one or more "reports" in a format known as a "workbook."

Generally, to generate a workbook request, the user logs into the server system through username and password combinations. The user interface then presents the user with options for designing a new workbook or modifying an existing workbook. To add a new workbook or modify an existing one, the user identifies the workbook (e.g., by naming it), and then selects one or more predefined reports or template/filter combinations to be included in the workbook. The user also selects the sequence of the reports and template/filter combinations. The server system may present the

available reports and template/filter combinations by accessing the data warehouse, as for example, with a ROLAP engine by displaying the available table entries from the relational database structure in the data warehouse. Auto-prompted and drill-down reports may also be included in a workbook.

5           A workbook is processed as a collection of reports, with each report being processed and then when all reports in the workbook are complete, the workbook is formulated and the users requesting the workbook are then notified. The reports in a workbook may be processed asynchronously or sequentially depending on the types of reports.

10           According to an embodiment of the present invention, the user may specify the format of each report in the workbook or to a stand-alone report being requested. As such, predefined or user defined formats may be selected and have that format map to many formatting settings. For example, the user may pick the "elegant" auto-format, which may be defined by a combination of fonts, styles, colors, banding, spacing, etc.

15           The entire report or workbook for which that format is selected may then have the "elegant" auto-format applied. In this embodiment the system may enable the user to select between multiple levels of formats that are applied in a hierarchical fashion. For example, four levels of formats may be specified as follows.

          The highest level format may be a global format that applies to all reports in

20           the workbook. Global reporting provides the base level of formatting or a default for all reports that are not otherwise specified. Global formatting may be over-ridden by lower levels of formats. This formatting applied may be, for example, formatting used in GridX AutoFormat. Formats that may be selected as global formats include

an AutoFormat style that applies a format to all reports in the workbook, and server or client macros that run when the workbook is executed by the OLAP system. Server macros allow the user to manipulate the workbook report data before it is returned to the browser on the user system. After being run once on the OLAP server, the workbook may be cached for all users of the system to access and manipulate the workbook through the use of server macros. Client macros may be launched to automatically manipulate the display of data already returned to the user interface system, such as a browser, to make additional functionality available. Global formats may also include provision of a worksheet for performance and error handling including debug mode and log statistics.

The next lower level format may be a report level format. In this level, the formats specified apply only to a particular report and not to other reports in the workbook. This format may position a report within a workbook or worksheet and select an AutoFormat for a report.

Autoformat comprises predefined format styles by specifying parameters for an ActiveX type application, for example. The parameters include name, gridline, font, banding. Gridline defines which of the following gridline properties apply to the report: topbar, bottombar, leftbar, rightbar, topsep, leftsep, totalssep, hgrid, and vgrid. Font defines which of the following font properties apply: genfont, toplabel, leftlabel, valuefont, totalslabel, totalsvalue, leveltop, levelleft, and levelvalue. Banding defines which of the following banding properties are applied: level, width, darkcolor, lightcolor, and inlabel.

The topbar gridline property determines the content of the top gridline of the report. The bottombar gridline determines the content of the bottom gridline of the report. The leftbar gridline determines the content of the leftmost gridline of the report. the rightbar gridline determines the content of the rightmost gridline of the report. The topsep determines the content of the line that separates the title labels in the top header area of the report from the raw data. The leftsep determines the line that separates the title labels in the left header area of a report from the raw data. The totalsset determines the line that separates totals from other data values. The Hgrid determines how horizontal lines should be applied to the data area of the report. The Vgrid determines how vertical lines should be applied to the data area of a report.

The genfont attribute determines the font style of all cells in a report unless the cells are defined by one or more of the other font properties. The toplabel font is the font for the top label area (column headers). The leftlabel font is the font for the left label area (row headers). The valuefont is the font of the data value area. The totalslabel is the value of the label of the totals area of the report. The totalsvalue is the font for the totals value of the report. The leveltop is the font for the top labels at a specific level. The levelleft font defines the font for the left labels at a specific level. The levelvalue is the font of the data values at a specific row level.

For banding, the level defines the level to which banding should be applied. The width defines the number of banding rows or columns that should be displayed with the same color. The darkcolor defines the dark band background color. The lightcolor defines the light band background color. The inlabel defines whether or not banding extends into the headers in the label areas.

Through selecting values for these attributes, the autoformat applies to define the format of a report in the workbook.

Autoprompt reports are reports that prompt the user to input some of the parameters, such as the filter criteria, for example, for the report when it is initiated.

5       Positioning a report within a worksheet may include vertical stacking which specifies the distance (in rows, for example) that the selected report may be placed from the report preceding it and the distance (in columns, for example) from the left sheet margin. If no report is specified above the selected report, then the system uses the top sheet margin as the bottom edge from which to count rows. Also, horizontal  
10       stacking may be provided which enables a user to specify the distance (in columns, for example) that the selected report should be placed from the report to the left of it and the distance (in rows, for example) from the top sheet margin. If no report is specified to the left of the selected report, then the system uses the left sheet margin as the edge from which to count columns. The system may also enable specification of the  
15       absolute cell reference from which the top left corner of the report is to be placed.

The next lower level format may be an attribute/element/metric level format which overrides the global and report level formats for certain attributes, elements, and metrics within a particular report. The next lower level format may comprise the value level format that overrides the global, report, and attribute/element/metric level  
20       format for a particular intersection of a metric and an element which establish a value.

Another level of formatting, a worksheet level formatting may also be applied. Worksheets may comprise collections of reports within the workbook and a worksheet

level format may be provided between the global and report level formats described above.

These formatting options for the workbooks may be selected by the user through a graphical user interface in the network user interface, such as the web browser. This process may be part of the request for a report to be processed by the OLAP system. Input may be through use of known techniques including direct entry into a field, pull-down menus, button selection, or other such techniques. Further, the formats may be from predefined styles that provide the parameters of the format. The user may also be able to define styles through a graphical user interface by saving formats input for a workbook or report as a style for later use with other workbooks or reports.

Also, formatting may be provided by the user through scripting. The user may use scripting, such as through the use of VBA or other scripting systems, to do completely programmatic and specialized formatting. This is possible because the data hooks (or "metadata") related to how formatting is applied are provided as parameters that may be used in scripts to modify formats. Through scripting, general and maintainable formatting may be provided rather than requiring hard-coded formatting within a grid or graph generated as a report.

Workbooks may be implemented through leveraging spreadsheet application programs, such as Microsoft Excel™. Through these spreadsheet applications, formatting of multiple reports into a combined workbook may be provided. Accordingly, the programming language in the spreadsheet application may be used to configure the format and arrangement of reports within the workbook. For example,

Visual Basic for Applications may be used in the Excel™ spreadsheet software application. This provides flexibility in formatting and combining reports to create new ones and in building an interactive spreadsheet application for the end user. For example, a spreadsheet application for a stock market workbook may provide pull-down menus for selecting different time ranges and industries which then update  
5 instantly/dynamically. Because these programming languages and spreadsheet programs, such as VBA and Excel are open standards, they are widely used and therefore, the usefulness of the workbook feature is enhanced.

Also, using Excel™, for example, enables users to select various properties  
10 and formatting for the reports that are to be generated. Integration of Excel™ with the invention allows users to use all of the functionality offered in Excel™ and to apply that functionality to the reports generated using the invention. For example, Excel™ offers existing macros that may be used to automate work flow or other features for the reports. Additionally, these macros may be modified or additional macros may be  
15 recorded to perform a specific operations desired by a user. The macros may be used, for example, to facilitate work flow. Generating reports in Excel™ also enables reports to display grids and graphs of the same report at the same time, display headers and footers, deliver multiple worksheets in a single workbook, support logical exceptions through use of overwriting rules, allows administrators to create formatted  
20 reports, etc.

A detailed embodiment of a system for enabling users to asynchronously generate report requests is depicted in Fig. 5. Such a system 50 comprises a data warehouse 12, server system 14, network output module 22, user systems 26, user

interface modules 28 and network 36. Additionally, server system 14 may be connected to a server cache 44 and network output module 22 may be connected to a network output module cache 46. Further, a network server 42 may be provided to generate output from network output module 22 across network 36 to user systems 26.

5 In system 50, network 36 may comprise an intranet, the Internet or any other network system that enables communication between a server and user system such as servers 14 and 42 and user systems 26. Network server 42 may comprise a web server and user system 26 may comprise computers connected over the Internet and World Wide Web to web server 42.

10 Server cache 44 may be provided to store calculated report results in a relational database structure. By providing a server cache 44 with server system 14, complex business reports may be accessed with one simple server request statement rather than having to access data warehouse 12 each time that a complex business report is requested. Additionally, through the use of server cache 44, load balancing  
15 may be provided. Further, for security purposes, server cache 44 may be used to index result tables by filter/template combination and by user.

Network output module cache 46 may be provided to store dynamic web pages created to display report results in files through network server 42. Cache results in network output module 46 may retain links to data warehouse content for *ad hoc*  
20 reporting and drilling.

Various security level architectures may be employed, although not depicted in Fig. 5. For example, a database level security may be provided between data warehouse 12 and server system 14. A network security user validation module may



be provided between network output module 22 and network server 42. Further, firewall and data encryption techniques may be provided within network 36 to prevent access to network server 42 from unqualified or unauthorized user systems 26. Also, an API may be provided to enable creation of further security level clearances to prevent user systems 26 access to data and data warehouse 12 for server cache 44 to which they do not have authorization.

Operation of system 50 may now be described. System 50 may enable users to request reports from server system 14 to be processed against data warehouse 12 through the use of a user interface module, such as a web browser operatively connected over the Internet to a network output module 22. As such a report may be requested by a user through user interface module 28 of user system 26.

To access the data warehouse, a user may first log in to network server 42 via user system 26. The user may do so by selecting a particular web site maintained by network server 42 or some other mechanism may be used. Network server 42 then presents the user with a view using user interface module 28 (*e.g.*, a web browser) comprising one or more options. A view may comprise a page or series of pages from a website, for example. One option presented in the view may be a mechanism that enables a user to generate a report request through this mechanism in the view. User interface module 28 may be used to submit a request for a report selected by the user to network server 42 over network 36. Network server 42 communicates the request to network output module 22. Network output module 22 receives the request and formulates a report request to be sent to server system 14 for processing using data warehouse 12. Network output module 22 may also perform other steps, as described

in detail below with respect to Fig. 3. Network output module 22 may also cooperate with an agent module operatively connected over system 50 to formulate a report. Server system 14 then accesses data warehouse 12 to perform all data retrieval and calculations to provide the report.

5           Network output module 22 may provide the functionality to determine whether a report request is the same or substantially the same as a previously requested report prior to forwarding a report to server system 14 for processing. Upon receipt of a report request, network output module 22 also controls network server 42 to present a new view to the user through user interface module 28 via network 36 while  
10           continuing to monitor progress of submitted reports identified for the particular user. This enables users to perform other tasks and thereby provides the asynchronous nature of the present invention. For example, the user may operate other views, pages, sites or any other action using user interface module 28, without having to wait for the report generation to be complete before being able to do so.

15           After the report generation is complete, server system 14 communicates the report to network output module 22. Network output module 22 prepares the report for presenting to the user(s) that requested the report using user interface module 28. As such, network output module 22 controls asynchronous processing of reports requested by the users.

20           Fig. 6 illustrates various components and objects of system 50 that perform tasks and functions described above. These objects may reside anywhere within system 50, but the location of each according to one embodiment is described as well. System 50 comprises a report requesting object 112, a report request receiving object

114, a report request submitting object 116, an option presenting object 118, a selected option receiving object 120, and a selected option submitting object 122. Report requesting object 112 communicates with report request receiving object 114 to enable a user to request a report to be generated. Report requesting object 112 may  
5 comprise functionality provided in a web page presented through user interface module 28. Report request receiving object 114 may comprise a module that is a portion of network output module 22. Report request submitting object 116 submits the user's report request for generation. Report request submitting object 116 may comprise an object that is a portion of network output module 22.

10       After a request has been submitted, option presenting object 118 may present various options for the report. Selected option receiving object 120 communicates with option presenting object 118 to select the option(s) selected by the user. Options may include specification of the name, content, project, filter, template and other components of a report as may be used by server system 14. Selected option  
15 submitting object 122 then submits the options selected by the user for processing. A report generating object 124 generates the requested report asynchronously based on the options selected by the system user. A report list creating object 126 creates a report which lists one or more reports requested. If a request for a report that has already been requested, network output module 22 adds the request, including an  
20 identification of the system user requesting the report, to the report entry without generating another report. A status page creating object 128 creates a status page which identifies the status of one or more requested reports. The status page may include options which may be selected by the system user (e.g., cancel request).

Option presenting object 118 and selected option submitting object 122 may comprise functionality presented in a page to the user via user interface module 28. Other objects described may comprise components of network output module 22.

Through an administrator module within network output module 22, an administrator may also monitor performances of each report, filter reports by pending, error and completed requests, cancel or modify reports, and may resubmit a report if it creates an error. The administrator module may have restricted access to prevent users from canceling other user requests. The administrator module may also allow for setting global settings for the system. Further, the administrator module may also enable administrators to view history lists, status lists, and cancellation pages from the viewpoint of users on the system. This enables administrators to be able to see what the user is seeing and therefore, troubleshoot or answer questions about the performance and operation of the system.

As described above, therefore, the present invention provides the ability for users and system administrators to cancel report requests. If a user desires to cancel a report request, the user may submit a report request cancellation by, for example, clicking on a "cancel" button using a conventional computer mouse or by selecting "cancel" from a standard pull-down menu. After a report request cancellation has been submitted, the system removes the request for that user from the report entry in the report list and then determines whether any other requests by other users remain for that report entry. If the report entry contains additional requests for that report, the system continues generating the report. This method enables users to cancel respective report requests while continuing report generation for other system users.

Additionally, this reduces system overhead by generating a single report for multiple requests. For example, if several users request the same report, only one report will be generated unlike prior systems where reports are generated for each request submitted by a user. For example, if ten users submitted a single request for the same report, ten  
5 reports would be generated. The invention, however, generates only one report that is used for each of the requests submitted by the ten (10) system users.

Another feature of the invention enables notification to a system user of a completed report. Once a report is generated, the system notifies the user that requested the report via e-mail, telephone, fax, pager, or other method that the report  
10 is completed. This may be done using known broadcast technology. For example, the invention may use Microstrategy's DSS Broadcaster™ to implement this notification feature.

Users may request and format reports generated using standard spreadsheet software using a network interface device (*e.g.*, a web browser). The web browser  
15 may be, for example, Microsoft Internet Explorer™, Netscape Navigator™, etc. Use of a web browser allows users to easily manipulate and generate reports without requiring the user to have a specific application loaded on the user system.

The present invention may be used in an overall system comprising a plurality of other systems for OLAP and decision support. According to one embodiment of the  
20 invention, a decision support system 10 may be provided as depicted in Fig. 7. System 10 may comprise a data warehouse 12, a server system 14, an architect module 16, an administrator module 18, a broadcast module 20, a network output module 22, a plurality of user systems 26, and an object creation module 24. Data warehouse 12

may comprise any data warehouse or data mart as is known in the art, including a relational database management system ("RDBMS") or a multidimensional database management system ("MDBMS") or a hybrid of the two. As discussed above, data warehouse 12 consolidates source system transaction data, organizes that data by business subject areas called dimensions and often includes data history for long term periods, such as two to five years. A data mart includes dimensions which tend toward application-specific subjects instead of integrating the entire business in one database such as departmental analysis data, data mining, and mass mailing lists.

Server system 14 comprises an OLAP server system for accessing and managing data stored in data warehouse 12. Server system 14 may comprise a ROLAP engine, MOLAP engine or a HOLAP engine according to different embodiments. Specifically, server system 14 may comprise a multithreaded server for performing sophisticated analysis directly against data warehouse 12. One embodiment of server system 14 may comprise a ROLAP engine known as DSS Server™ offered by MicroStrategy.

Architect module 16 may comprise a module that enables developers to create and maintain data and metadata in data warehouse 12. Metadata may be considered to be data about data, such as data element descriptions, data type descriptions, attributes/property descriptions, range/domain descriptions, and process/method descriptions. Data and metadata stored in data warehouse 12 may thus be modified and organized by architect module 16. According to one embodiment of the invention, architect module 16 may comprise a software package known as DSS Architect™ offered by MicroStrategy.

Administrator module 18 may comprise a module for facilitating the development, deployment and management of data warehouse applications supporting large volumes of users over various distribution mechanisms. Administrator module 18 may comprise an object manager and a warehouse monitor. The object manager  
5 allows objects to be shared across databases for easy migration from development to production. The warehouse monitor provides performance monitoring and management tools to support thousands of users across a distributive database environment. The warehouse monitor collects statistics for the purpose of identifying performance bottlenecks, warehouse tuning, and cost analysis. According to one  
10 embodiment of the invention, administrator module 18 may comprise a module known as DSS Administrator™ offered by MicroStrategy.

Server system 14 also connects to an object creation module 24. Object creation module 24 may comprise a component object model ("COM") application program interface ("API") for custom decision support development. According to  
15 one embodiment of the invention, object creation module 24 may comprise a software module known as DSS Objects™ offered by MicroStrategy. Additionally, custom applications may interface with object creation module 24 including Delphi, Visual Basic, and C++ programming modules.

Server system 14 also interfaces with a number of devices that enable users to  
20 access the contents of the data warehouse to perform various data analyses. One embodiment of such a system comprises a user system 26 with an agent module 28 operating thereon. Agent module 28 may comprise a decision support interface to

enable a user to provide integrated query and reporting, powerful analytics, and decision support through any standard user system 26.

Agent module 28 may operate on any user system including personal computers, network workstations, laptop computers or any other electronic device  
5 connected to server system 14. Agent module 28 may enable the user to define queries to be performed against the data contained in data warehouse 12 using components, templates, filters, reports, agents, etc. Components may include dimensions, attributes, attribute elements, and metrics; in other words, the building blocks for templates, filters, and reports. Templates generally define a report format  
10 and specify the attributes, dimensions, metrics, and display properties comprising a report. Filters generally qualify report content and identify the subset of data warehouse 12 to be included in a report. Reports are generally understood to be a data analysis created by combining a template (the format) with a filter (the content). Agents are a group of reports cached on a time -or event- based schedule for rapid  
15 retrieval and batch processing.

Accordingly, agent module 28 enables a user access to the contents of data warehouse 12 to provide detailed analysis on an *ad hoc* basis. According to one embodiment of the invention, agent module 28 may comprise a software package known as DSS Agent™ offered by MicroStrategy. One of the advantages of DSS  
20 Agent™ includes its use of a ROLAP architecture on server system 14 and a RDBMS in data warehouse 12 to provide a more scaleable environment. Through DSS Agent™, a user can “drill down” which allows the user to dynamically change the level of detail in a report. Drilling down allows the user to drill to a lower level



attribute so that the resulting report displays data with a greater level of detail. For example, one can drill down from year to month to week to day. DSS Agent™ also enables users to “drill up” to a higher level attribute. Drilling up summarizes the selected data to a higher level total. For example, one can drill from day to week to  
5 month to year. DSS Agent™ also enables a user to “drill within.” Drilling within allows a user to go to a different hierarchy within the same dimension. Drilling within is often used to examine the characteristics of selected data. For example, drilling within enables a user to drill from item to color when looking at a particular retail item such as an automobile, clothing or the like. Drilling across allows the user to  
10 drill to an altogether different dimension or subject area. For example, one can drill across from a region to a month. Accordingly, through use of agent module 28, server system 14, and data warehouse 12, drilling is a powerful tool that is easily implemented using a ROLAP architecture which is not as easily accessible in MOLAP unless a user also implements the complicated structure of a HOLAP architecture.

15       Along with agent module 28, user system 26 may also include a report writing module 30, an executive module 32, and a spreadsheet module 34. Report writing module 26 may comprise an OLAP report writer. Executive module 32 may comprise a module design tool for developing custom EIS applications. This module is a design tool for developing briefing books that provide high level users with a series of  
20 views that describe their business. Once created, end users can access briefing books through agent module 28 in EIS mode. Such a system is easily implemented with agent module 28 by compiling sets of analyses into dynamic pages that immediately

focus users on their key business drivers. One embodiment of executive module 32 may comprise the software made DSS Executive™ offered by MicroStrategy.

Spreadsheet module 34 may comprise an add on to existing spreadsheet programs or may comprise an entirely new spreadsheet program. Spreadsheet module  
5 34 may enable reports and analyses generated from agent module 28 to be presented in a traditional spreadsheet program format to enable users to view results in preexisting front end interfaces. Spreadsheet module 34 may comprise the Microsoft Excel™ spreadsheet program offered by Microsoft and an Excel™ Add-In program offered by MicroStrategy.

10 Another module for accessing content of server system 14 may comprise a network output module 22. Network output module 22 offers the advantage of enabling user system 26 access to server system 14 in data warehouse 12 without requiring an additional agent module 28 to be stored on user system 26. Instead, user system 26 may have a user interface module 38 residing on user system 26. User  
15 interface module 38 may comprise any module that enables a user system to interface with the network output module over a network 36. According to one embodiment of the invention, network 36 may comprise an intranet or the Internet or other developed Internet-type networks. Further, user interface module 38 may comprise any standard browser module such as Microsoft Internet Explorer™ or Netscape Navigator™. As  
20 most user systems 26 already have a user interface module 38 stored and operating thereon, network output module 22 offers the advantage of enabling users access to server system 14 in data warehouse 12 without learning to operate a new module such as agent module 28. One embodiment of network output module 22 may comprise a

web-based module called DSS Web™ offered by MicroStrategy. Accordingly, in one embodiment, a user can access server system 14 through a standard web browser, such as Microsoft Internet Explorer™, over the Internet through network output module 22, such as DSS Web™.

5 In this embodiment, network output module 22 may comprise a World Wide Web tool used in conjunction with server system 14 for allowing users to deploy data warehouse/decision support applications over the Internet using industry standard World Wide Web browsers as a client. As a result, a user can access the data warehouse with little or no client maintenance, little or no software to install, and only  
10 a small amount of additional training while still maintaining all of the capabilities of agent module 28. One embodiment of such a network output module 22 comprises DSS Web™ offered by MicroStrategy. Such an embodiment provides a broad array of options for viewing information sets, such as spreadsheet grids and a wide variety of graphs. Through this module's exceptional reporting capabilities, users receive key  
15 elements of a report in easily interpretable, plain messages. This module also allows users to "drill" dynamically to a lower level of detail to view the underlying information or to create and save new analyses. For sensitive information, this module provides security plug-ins that allow the user to extend the standard security functionality with additional user authentication routines. This module may also  
20 provide an interface API that allows users to customize, integrate, and imbed this functionality into other applications. For example, the data syndicator for health care information could utilize this module with a customize interface to sell access to

health care information to Health Maintenance Organizations, hospitals, and pharmacies.

Another module that enables user access to the data warehouse 12 comprises a broadcast module 20. Broadcast module 20 may comprise an information broadcast server designed to deliver personalized messages to multiple recipients via user devices 40. User devices 40 may comprise electronic mail, facsimile, pager, mobile phone, telephone, and multiple other types of user information devices. Broadcast module 20, like agent module 28 and network output module 22, enables users to define queries and reports that are to be run against server system 14 and data warehouse 12. Broadcast module 20 may then send personalized information to users at predefined intervals or when criteria specified in their reports exceed predefined thresholds. Broadcast module 20 may continually monitor and run such reports to enable users to receive appropriate information as it becomes available. Broadcast module 20 may also personalize the report to ensure that users see only that portion of a report that is relative to that user and thereby frees analysts from acting as the information provider. Broadcast module 20 may also provide a wide range of content formatting options to enable users to receive information in a form that they can best understand. One such formatting may include formatting a request for a pager, facsimile, mobile phone, electronic mail, web or other output user device 40.

Broadcast module 20 may specify a mechanism called a service for generating information and a subscription that determines which users receive information from which service. Accordingly, a service may be tied to, for example, a specific report, series of reports or elements within a report, etc. Also, subscribers may include users,

groups of users or only specific user devices 40 for a particular user. Broadcast module 20 may output information to user devices 40 using push technology. In other words, as soon as a criteria is established, broadcast module 20 generates a report to a user device 40. Accordingly, a user can thus have up-to-date information about the contents of data warehouse 12 without having to submit a query or log-in to a software module such as agent module 28 or network output module 22 to initiate the query.

In such an overall system, the present invention may enable asynchronous report generation through operation of network output module system 22 as described above. Other embodiments and uses of the invention will be apparent to those skilled in the art from consideration of the specification and practice of the invention disclosed herein. The specification and examples should be considered exemplary only. The scope of the invention is only limited by the claims appended hereto.